




on each node). These operations are illustrated in Fig. 2. The cost of elementary operations is given by two functions,  $c_{\text{del}}$  and  $c_{\text{rel}}$ , where  $c_{\text{del}}(\tau)$  is the cost of deleting or inserting a vertex with label  $\tau$ , and  $c_{\text{rel}}(\tau_1, \tau_2)$  is the cost of changing the label of a vertex from  $\tau_1$  to  $\tau_2$ . A deletion in  $F$  is equivalent to an insertion in  $G$  and vice versa, so we can focus on finding the minimum cost of a sequence of deletions and relabels in both trees that transform  $F$  and  $G$  into isomorphic trees.



**Fig. 2.** The three editing operations on a tree with vertex labels.

*Previous results.* To state running times, we need some basic notation. Let  $n$  and  $m$  denote the sizes  $|F|$  and  $|G|$  of the two input trees, ordered so that  $n \geq m$ . Let  $n_{\text{leaves}}$  and  $m_{\text{leaves}}$  denote the corresponding number of leaves in each tree, and let  $n_{\text{height}}$  and  $m_{\text{height}}$  denote the corresponding height of each tree, which can be as large as  $n$  and  $m$  respectively.

Tai [10] presented the first algorithm for computing tree edit distance, which requires  $O(n_{\text{leaves}}^2 m_{\text{leaves}}^2 nm)$  time and space. Tai's algorithm thus has a worst-case running time of  $O(n^3 m^3) = O(n^6)$ . Shasha and Zhang [9] improved this result to an  $O(\min\{n_{\text{height}}, n_{\text{leaves}}\} \cdot \min\{m_{\text{height}}, m_{\text{leaves}}\} \cdot nm)$  time algorithm using  $O(nm)$  space. In the worst case, their algorithm runs in  $O(n^2 m^2) = O(n^4)$  time. Klein [6] improved this result to a worst-case  $O(m^2 n \log n) = O(n^3 \log n)$ -time algorithm that

uses  $O(nm)$  space. In addition, Klein's algorithm can be adapted to solve an unrooted version of the problem. These last two algorithms are based on closely related dynamic programs, and both present different ways of computing only a subset of a larger dynamic program table; these entries are referred to as *relevant subproblems*. In [4], Dulucq and Touzet introduced the notion of a *decomposition strategy* (see Section 2.3) as a general framework for algorithms that use this type of dynamic program, and proved a lower bound of  $\Omega(nm \log n \log m)$  time for any such strategy.

Many other solutions have been developed; see [1, 11] for surveys. The most recent development is by Chen [3], who presented a different approach that uses results on fast matrix multiplication. Chen's algorithm uses  $O(nm + nm_{\text{leaves}}^2 + n_{\text{leaves}} m_{\text{leaves}}^{2.5})$  time and  $O(n + (m + n_{\text{leaves}}^2) \min\{n_{\text{leaves}}, n_{\text{height}}\})$  space. In the worst case, this algorithm runs in  $O(nm^{2.5}) = O(n^{3.5})$  time. In general, Klein's algorithm remained the best in terms of worst-case time complexity.

*Our results.* In this paper, we present a new algorithm for tree edit distance that falls into the same *decomposition strategy* framework of [6, 9, 4]. Our algorithm runs in  $O(nm^2(1 + \log \frac{n}{m})) = O(n^3)$  worst-case time and  $O(nm)$  space, and can be adapted for the case where the trees are not rooted. The corresponding edit script can easily be obtained within the same time and space bounds. We therefore improve upon all known algorithms in the worst-case time complexity. Our approach is based on Klein's, but whereas the recursion scheme in Klein's algorithm is determined by just one of the two input trees, in our algorithm the recursion depends alternately on both trees. Furthermore, we prove a worst-case lower bound of  $\Omega(nm^2(1 + \log \frac{n}{m}))$  time on all decomposition strategy algorithms. This bound improves the previous best lower bound of  $\Omega(nm \log n \log m)$  time [4], and establishes the optimality of our algorithm among all decomposition strategy algorithms. Our algorithm is simple, making it easy to implement, but both the upper and lower bound proofs require complicated analysis.

*Roadmap.* In Section 2 we give simple and unified presentations of the two well-known tree edit algorithms, on which our algorithm is based, and the class of decomposition strategy algorithms. We present and analyze our algorithm in Section 3, and prove the matching lower bound in Section 4. We conclude in section 5.

## 2 Background and Framework

Both the existing algorithms and ours compute the edit distance of finite ordered  $\Sigma$ -labeled forests, henceforth *forests*. The unique empty forest/tree is denoted by  $\emptyset$ . The vertex set of a forest  $F$  is written simply as  $F$ , as when we speak of a vertex  $v \in F$ . For any forest  $F$  and  $v \in F$ ,  $\sigma(v)$  denotes the  $\Sigma$ -label of  $v$ ,  $F_v$  denotes the subtree of  $F$  rooted at  $v$ , and  $F - v$  denotes the forest obtained from  $F$  after deleting  $v$ . The leftmost and rightmost trees of  $F$  are denoted by  $L_F$  and  $R_F$  and their roots by  $\ell_F$  and  $r_F$ . We denote by  $F - L_F$  the forest obtained from  $F$  after deleting the entire leftmost tree  $L_F$ ; similarly  $F - R_F$ . A forest obtained from  $F$  by a sequence of any number of deletions of the leftmost and rightmost roots is called a *subforest* of  $F$ .

Given forests  $F$  and  $G$  and vertices  $v \in F$  and  $w \in G$ , we write  $c_{\text{del}}(v)$  instead of  $c_{\text{del}}(\sigma(v))$  for the cost of deleting or inserting  $v$ , and we write  $c_{\text{rel}}(v, w)$  instead of  $c_{\text{rel}}(\sigma(v), \sigma(w))$  for the cost relabeling  $v$  to  $w$ .  $\delta(F, G)$  denotes the edit distance between the forests  $F$  and  $G$ .

Because insertion and deletion costs are the same (for a node of a given label), *insertion in one forest* is tantamount to *deletion in the other forest*. Therefore, the only edit operations we need to consider are relabels and deletions of nodes in both forests. In the next two sections, we briefly present the algorithms of Shasha and Zhang, and of Klein. Our presentation is inspired by the tree similarity survey of Bille [1], and is essential for understanding our algorithm.

### 2.1 Shasha and Zhang's Algorithm [9]

Given two forests  $F$  and  $G$  of sizes  $n$  and  $m$  respectively, the following lemma is easy to verify.

Intuitively, this lemma says that the two rightmost roots in  $F$  and  $G$  are either matched with each other or one of them is deleted.

**Lemma 1 ([9]).**  $\delta(F, G)$  can be computed as follows:

- $\delta(\emptyset, \emptyset) = 0$
- $\delta(F, \emptyset) = \delta(F - r_F, \emptyset) + c_{\text{del}}(r_F)$
- $\delta(\emptyset, G) = \delta(\emptyset, G - r_G) + c_{\text{del}}(r_G)$
- $\delta(F, G) = \min \begin{cases} \delta(F - r_F, G) + c_{\text{del}}(r_F), \\ \delta(F, G - r_G) + c_{\text{del}}(r_G), \\ \delta(R_F - r_F, R_G - r_G) \\ \quad + \delta(F - R_F, G - R_G) \\ \quad + c_{\text{rel}}(r_F, r_G) \end{cases}$

The above lemma yields an  $O(m^2n^2)$  dynamic program algorithm: If we index the vertices of the forests  $F$  and  $G$  according to their postorder traversal position, then entries in the dynamic program table correspond to pairs  $(F', G')$  of subforests  $F'$  of  $F$  and  $G'$  of  $G$  where  $F'$  contains vertices  $\{i_1, \dots, j_1\}$  and  $G'$  contains vertices  $\{i_2, \dots, j_2\}$  for some  $1 \leq i_1 \leq j_1 \leq n$  and  $1 \leq i_2 \leq j_2 \leq m$ .

However, as we will presently see, only  $O(\min\{n_{\text{height}}, n_{\text{leaves}}\} \cdot \min\{m_{\text{height}}, m_{\text{leaves}}\} \cdot nm)$  different *relevant subproblems* are encountered by the recursion computing  $\delta(F, G)$ . We calculate the number of *relevant subforests* of  $F$  and  $G$  independently, where a forest  $F'$  (respectively  $G'$ ) is a relevant subforest of  $F$  (respectively  $G$ ) if it shows up in the computation of  $\delta(F, G)$ . Clearly, multiplying the number of relevant subforests of  $F$  and of  $G$  is an upper bound on the total number of relevant subproblems.

We focus on counting the number of relevant subforests of  $F$ . The count for  $G$  is similar. First, notice that for every node  $v \in F$ ,  $F_v - v$  is a relevant subproblem. This is because the recursion allows us to delete the rightmost root of  $F$  repeatedly until  $v$  becomes the rightmost root; we then match  $v$  (i.e., relabel it) and get the desired relevant subforest. A more general claim is stated and proved later on in Lemma 3. We define  $\text{keyroots}(F) = \{\text{the root of } F\} \cup \{v \in F \mid v \text{ has a left sibling}\}$ . Every relevant subforest of

$F$  is a prefix (with respect to the postorder indices) of  $F_v - v$  for some node  $v \in \text{keyroots}(F)$ . If we define  $\text{cdepth}(v)$  to be the number of keyroot ancestors of  $v$ , and  $\text{cdepth}(F)$  to be the maximum  $\text{cdepth}(v)$  over all nodes  $v \in F$ , we get that the total number of relevant subforest of  $F$  is at most

$$\begin{aligned} \sum_{v \in \text{keyroots}(F)} |F_v| &= \sum_{v \in F} \text{cdepth}(v) \\ &\leq \sum_{v \in F} \text{cdepth}(F) \\ &= |F| \text{cdepth}(F). \end{aligned}$$

This means that given two trees,  $F$  and  $G$ , of sizes  $n$  and  $m$  we can compute  $\delta(F, G)$  in  $O(\text{cdepth}(F)\text{cdepth}(G)mn)$  time. Shasha and Zhang also proved that for any tree  $T$  of size  $n$ ,  $\text{cdepth}(T) \leq \min\{n_{\text{height}}, n_{\text{leaves}}\}$ , hence the result. In the worst case, this algorithm runs in  $O(m^2n^2) = O(n^4)$  time.

## 2.2 Klein's Algorithm [6]

Klein's algorithm is based on a recursion similar to Lemma 1. Again, we consider forests  $F$  and  $G$  of sizes  $|F| = n \geq |G| = m$ . Now, however, instead of recursing always on the rightmost roots of  $F$  and  $G$ , we recurse on the leftmost roots if  $|L_F| \leq |R_F|$  and on the rightmost roots otherwise. In other words, the "direction" of the recursion is determined by the (initially) larger of the two forests. We assume the number of relevant subforests of  $G$  is  $O(m^2)$ ; we have already established that this is an upper bound.

We next show that Klein's algorithm yields only  $O(n \log n)$  relevant subforests of  $F$ . The analysis is based on a technique called *heavy path decomposition* introduced by Harel and Tarjan [5]. Briefly: we mark the root of  $F$  as *light*. For each internal node  $v \in F$ , we pick one of  $v$ 's children of maximum size and mark it as *heavy*, and we mark all the other children of  $v$  as *light*. We define  $\text{ldepth}(v)$  to be the number of light nodes that are ancestors of  $v$  in  $F$ , and  $\text{light}(F)$  as the set of all light nodes in  $F$ . By [5], for any forest  $F$  and vertex  $v \in F$ ,  $\text{ldepth}(v) \leq \log |F| + O(1)$ .

Note that every relevant subforest of  $F$  is obtained by some  $i \leq |F_v|$  many consecutive deletions from  $F_v$  for some light node  $v$ . Therefore, the total number of relevant subforests of  $F$  is at most

$$\begin{aligned} \sum_{v \in \text{light}(F)} |F_v| &= \sum_{v \in F} \text{ldepth}(v) \\ &\leq \sum_{v \in F} (\log |F| + O(1)) \\ &= O(|F| \log |F|). \end{aligned}$$

Thus, we get an  $O(m^2n \log n) = O(n^3 \log n)$  algorithm for computing  $\delta(F, G)$ .

## 2.3 The Decomposition Strategy Framework

Both Klein's and Shasha and Zhang's algorithms are based on Lemma 1. The difference between them lies in the choice of when to recurse on the rightmost roots and when on the leftmost roots. The family of *decomposition strategy* algorithms based on this lemma was formalized by Dulucq and Touzet in [4].

**Definition 1 (Strategy).** *Let  $F$  and  $G$  be two forests. A strategy is a mapping from pairs  $(F', G')$  of subforests of  $F$  and  $G$  to  $\{\text{left}, \text{right}\}$ .*

Each strategy is associated with a specific set of recursive calls (or a dynamic program algorithm). The strategy of Shasha and Zhang's algorithm is  $S(F', G') = \text{right}$  for all  $F', G'$ . The strategy of Klein's algorithm is  $S(F', G') = \text{left}$  if  $|L_{F'}| \leq |R_{F'}|$ , and  $S(F', G') = \text{right}$  otherwise. Notice that Shasha and Zhang's strategy does not depend on the input trees, while Klein's strategy depends only on the larger input tree. Dulucq and Touzet proved a lower bound of  $\Omega(mn \log m \log n)$  time for any strategy based algorithm.


## 3 The Algorithm

In this section we present our algorithm for computing  $\delta(F, G)$  given two trees  $F$  and  $G$  of sizes  $|F| = n \geq |G| = m$ . The algorithm recursively uses Klein's strategy in a divide-and-conquer

manner to achieve  $O(nm^2(1 + \log \frac{n}{m})) = O(n^3)$  running time in the worst case. The algorithm's space complexity is  $O(nm)$ . We begin with the observation that Klein's strategy always determines the direction of the recursion according to the  $F$ -subforest, even in subproblems where the  $F$ -subforest is smaller than the  $G$ -subforest. However, it is not straightforward to change this since even if at some stage we decide to switch to Klein's strategy based on the other forest, we must still make sure that all subproblems previously encountered are entirely solved. At first glance this seems like a real obstacle since apparently we only add new subproblems to those that are already computed.

For clarity we describe the algorithm recursively. A dynamic programming description and a proof of the  $O(mn)$  space complexity will appear in the full version of this paper.

For a tree  $F$  of size  $n$ , define the set  $\text{TopLight}_F$  to be the set of roots of the forest obtained by removing the heavy path of  $F$  (i.e., the unique path starting from the root along heavy nodes). Note that  $\text{TopLight}_F$  is the set of light nodes with ldepth 1 in  $F$  (see the definition of ldepth in section 2.2). This definition is illustrated in Fig. 3. Note that the following two con-



**Fig. 3.** A tree  $F$  with  $n$  nodes. The black nodes belong to the heavy path. The white nodes are in  $\text{TopLight}_F$ , and the size of each subtree rooted at a white node is at most  $\frac{n}{2}$ .

ditions are always satisfied:

$$(*) \quad \sum_{v \in \text{TopLight}_F} |F_v| \leq n.$$

This follows from the fact that  $F_{v'}$  and  $F_{v''}$  are disjoint for any  $v', v'' \in \text{TopLight}_F$ .

$$(**) \quad |F_v| < \frac{n}{2} \text{ for every } v \in \text{TopLight}_F, \text{ since otherwise } v \text{ would be a heavy node.}$$

*The Algorithm.* We compute  $\delta(F, G)$  recursively as follows:

- (1) If  $|F| < |G|$ , compute  $\delta(G, F)$  instead. That is, we order the pair  $\{F, G\}$  such that  $F$  is always the larger forest.
- (2) Recursively compute  $\delta(F_v, G)$  for all  $v \in \text{TopLight}_F$ . Note that along the way this computes  $\delta(F_{v'} - v', G_w - w)$  for all  $v'$  not in the heavy path of  $F$  and for all  $w \in G$ .
- (3) Compute  $\delta(F, G)$  using Klein's strategy (matching and deleting either from the left or from the right according to the larger of  $F$  and  $G$ ). Do not recurse into subproblems that were previously computed in step (2).

The correctness of the algorithm follows immediately from the correctness of Klein's algorithm. The algorithm is evidently a decomposition strategy algorithm, since for all subproblems, it either deletes or matches the leftmost or rightmost roots.

*Time Complexity.* We show that our algorithm has a worst-case runtime of  $O(m^2n(1 + \log \frac{n}{m})) = O(n^3)$ .

We proceed by counting the number of subproblems computed in each step of the algorithm. Let  $R(F, G)$  denote the number of relevant subproblems encountered by the algorithm in the course of computing  $\delta(F, G)$ .

In step (2) we compute  $\delta(F_v, G)$  for all  $v \in \text{TopLight}_F$ . Hence, the number of subproblems encountered in this step is  $\sum_{v \in \text{TopLight}_F} R(F_v, G)$ .

In step (3) we compute  $\delta(F, G)$  using Klein's strategy. We bound the number of relevant subproblems by multiplying the number of relevant subforests in  $F$  and in  $G$ . For  $G$ , we count all possible  $O(|G|^2)$  subforests obtained by left and right deletions. Note that for any node  $v'$  not in

the heavy path of  $F$ , the subproblem obtained by matching  $v'$  with any node  $w$  in  $G$  was already computed in step (2). This is because any such  $v'$  is contained in  $F_v$  for some  $v \in \text{TopLight}_F$ , so  $\delta(F_{v'} - v', G_w - w)$  is computed in the course of computing  $\delta(F_v, G)$  (we prove this formally in Lemma 3). Furthermore, note that in Klein's algorithm, a node  $v$  on the heavy path of  $F$  cannot be matched or deleted until the remaining subforest of  $F$  is precisely  $F_v$ . At this point, both matching  $v$  or deleting  $v$  results in the same new relevant subforest  $F_v - v$ . This means that we do not have to consider matchings of nodes when counting the number of relevant subproblems in step (3). It suffices to consider only the  $|F|$  subforests obtained by deletions according to Klein's strategy. Thus, the total number of new subproblems encountered step (3) is bounded by  $|G|^2|F|$ .

We have established that  $R(F, G)$  is at most

$$|G|^2|F| + \sum_{v \in \text{TopLight}_F} R(F_v, G), \text{ if } |F| \geq |G|$$

$$|F|^2|G| + \sum_{w \in \text{TopLight}_G} R(F, G_w), \text{ if } |F| < |G|$$

We first show, by a crude estimate, that this leads to an  $O(n^3)$  runtime. Later, we analyze the dependency on  $m$  and  $n$  accurately.

**Lemma 2.**  $R(F, G) \leq 4(|F||G|)^{3/2}$ .

*Proof.* We proceed by induction on  $|F| + |G|$ . There are two symmetric cases. If  $|F| \geq |G|$  then  $R(F, G) \leq |G|^2|F| + \sum_{v \in \text{TopLight}_F} R(F_v, G)$ . Hence, by the inductive assumption,

$$\begin{aligned} R(F, G) &\leq |G|^2|F| + \sum_{v \in \text{TopLight}_F} 4(|F_v||G|)^{3/2} \\ &\leq |G|^2|F| + 4|G|^{3/2} \sum_{v \in \text{TopLight}_F} |F_v|^{3/2} \\ &\leq |G|^2|F| + 4|G|^{3/2} \sum_{v \in \text{TopLight}_F} |F_v| \max_{v \in \text{TopLight}_F} \sqrt{|F_v|} \\ &\leq |G|^2|F| + 4|G|^{3/2}|F|\sqrt{|F|/2} \\ &= |G|^2|F| + 2\sqrt{2}(|F||G|)^{3/2} \\ &\leq 4(|F||G|)^{3/2}. \end{aligned}$$

Here we have used facts (\*) and (\*\*) and the fact that  $|F| \geq |G|$ . The case where  $|F| < |G|$  is symmetric.  $\square$

This crude estimate gives a worst-case runtime of  $O(n^3)$ . We now analyze the dependence on  $m$  and  $n$  more accurately. Along the recursion defining the algorithm, we view step (2) as only making recursive calls, but not producing any relevant subproblems. Rather, every new relevant subproblem is created in step (3) for a unique recursive call of the algorithm. So when we count relevant subproblems, we sum the number of new relevant subproblems encountered in step (3) over all recursive calls to the algorithm.

We define sets  $A, B \subseteq F$  as follows:

$$A = \{a \in \text{light}(F) : |F_a| \geq m\}$$

$$B = \{b \in F - A : b \in \text{TopLight}_{F_a} \text{ for some } a \in A\}.$$

Note that the root of  $F$  belongs to  $A$ . We count separately:

- (i) the relevant subproblems created in just step (3) of recursive calls  $\delta(F_a, G)$  for all  $a \in A$ , and
- (ii) the relevant subproblems encountered in the entire computation of  $\delta(F_b, G)$  for all  $b \in B$  (i.e.,  $\sum_{b \in B} R(F_b, G)$ ).

Together, this counts all relevant subproblems for the original  $\delta(F, G)$ . To see this, consider the original call  $\delta(F, G)$ . Certainly, the root of  $F$  is in  $A$ . So all subproblems generated in step (3) of  $\delta(F, G)$  are counted in (i). Now consider the recursive calls made in step (2) of  $\delta(F, G)$ . These are precisely  $\delta(F_v, G)$  for  $v \in \text{TopLight}_F$ . For each  $v \in \text{TopLight}_F$ , notice that  $v$  is either in  $A$  or in  $B$ ; it is in  $A$  if  $|F_v| \geq m$ , and in  $B$  otherwise. If  $v$  is in  $B$ , then all subproblems arising in the entire computation of  $\delta(F_v, G)$  are counted in (ii). On the other hand, if  $v$  is in  $A$ , then we are in analogous situation with respect to  $\delta(F_v, G)$  as we were in when we considered  $\delta(F, G)$  (i.e., we count separately the subproblems created in step (3) of  $\delta(F_v, G)$  and the subproblems coming from  $\delta(F_u, G)$  for  $u \in \text{TopLight}_{F_v}$ ).

Earlier in this section, we saw that the number of subproblems created in step (3) of  $\delta(F, G)$  is  $|G|^2|F|$ . In fact, for any  $a \in A$ , by the same

argument, the number of subproblems created in step (3) of  $\delta(F_a, G)$  is  $|G|^2|F_a|$ . Therefore, the total number of relevant subproblems of type (i) is  $|G|^2 \sum_{a \in A} |F_a|$ . For  $v \in F$ , define  $\text{depth}_A(v)$  to be the number of ancestors of  $v$  that lie in the set  $A$ . We claim that  $\text{depth}_A(v) \leq 1 + \log \frac{n}{m}$  for all  $v \in F$ . To see this, consider any sequence  $a_0, \dots, a_k$  in  $A$  where  $a_i$  is a descendent of  $a_{i-1}$  for all  $i \in [1, k]$ . Note that  $|F_{a_i}| \leq \frac{1}{2}|F_{a_{i-1}}|$  for all  $i \in [1, k]$  since the  $a_i$  are light nodes, and note that  $|F_{a_k}| \geq m$  by the definition of  $A$ . It follows that  $k \leq \log \frac{n}{m}$ , i.e.,  $A$  contains no sequence of descendants of length  $> 1 + \log \frac{n}{m}$ . So clearly every  $v \in F$  has  $\text{depth}_A(v) \leq 1 + \log \frac{n}{m}$ .

We now have the number of relevant subproblems of type (i) as

$$\begin{aligned} |G|^2 \sum_{a \in A} |F_a| &= m^2 \sum_{v \in F} \text{depth}_A(v) \\ &\leq m^2 \sum_{v \in F} (1 + \log \frac{n}{m}) \\ &= m^2 n (1 + \log \frac{n}{m}). \end{aligned}$$

The relevant subproblems of type (ii) are counted by  $\sum_{b \in B} R(F_b, G)$ . Using Lemma 2, we have

$$\begin{aligned} \sum_{b \in B} R(F_b, G) &\leq 4|G|^{3/2} \sum_{b \in B} |F_b|^{3/2} \\ &\leq 4|G|^{3/2} \sum_{b \in B} |F_b| \max_{b \in B} \sqrt{|F_b|} \\ &\leq 4|G|^{3/2} |F| \sqrt{m} = 4m^2 n. \end{aligned}$$

Here we have used the facts that  $|F_b| < m$  and  $\sum_{b \in B} |F_b| \leq |F|$  (since the trees  $F_b$  are disjoint for different  $b \in B$ ). Therefore, the total number of relevant subproblems for  $\delta(F, G)$ —and hence the runtime of the algorithm—is at most  $m^2 n (1 + \log \frac{n}{m}) + 4m^2 n = O(m^2 n (1 + \log \frac{n}{m}))$ .

*Unrooted Trees.* Our algorithm can be adapted to compute edit distance of *unrooted ordered trees*. An unrooted ordered tree is an acyclic graph with a cyclic ordering defined on the edges incident on each node in the graph. In the modified algorithm, we arbitrarily choose a root for the larger of the two trees. We change the first recursive level of the algorithm, so that it now

computes the edit distance with respect to any possible choice of a root for the smaller tree. This does not change the time complexity since the number of different relevant subforests for a tree of size  $m$  is bounded by  $m^2$  whether we consider a single choice for the root or all possible choices. This idea will be described in detail in the full version of this paper.

## 4 A Tight Lower Bound for Strategy Algorithms

In this section we present a lower bound on the worst-case runtime of strategy algorithms. We first give a simple proof of an  $\Omega(m^2 n)$  lower bound. In the case where  $m = \Theta(n)$ , this gives a lower bound of  $\Omega(n^3)$  which shows that our algorithm is worst-case optimal among all strategy-based algorithms. To prove that our algorithm is worst-case optimal for any  $m \leq n$ , we analyze a more complicated scenario that gives a lower bound of  $\Omega(m^2 n (1 + \log \frac{n}{m}))$ , matching the running time of our algorithm.

In analyzing strategies we will use the notion of a *computational path*, which corresponds to a specific sequence of recursion calls. Recall that for all subforest-pairs  $(F', G')$ , the strategy  $S$  determines a direction: either *right* or *left*. The recursion can either delete from  $F'$  or from  $G'$  or match. A computational path is the sequence of operations taken according to the strategy in a specific sequence of recursive calls. For convenience, we sometimes describe a computational path by the sequence of subproblems it induces, and sometimes by the actual sequence of operations: either “delete from the  $F$ -subforest”, “delete from the  $G$ -subforest”, or “match”.

The following lemma states that every strategy computes the edit distance between every two root-deleted subtrees of  $F$  and  $G$ .

**Lemma 3.** *For any strategy  $S$ , the pair  $(F_v - v, G_w - w)$  is a relevant subproblem for all  $v \in F$  and  $w \in G$ .*

*Proof.* First note that a node  $v' \in F_v$  (respectively,  $w' \in G_w$ ) is never deleted or matched before  $v$  (respectively,  $w$ ) is deleted or matched. Consider the following computational path:

- Delete from  $F$  until  $v$  is either the leftmost or the rightmost root.
- Next, delete from  $G$  until  $w$  is either the leftmost or the rightmost root.

Let  $(F', G')$  denote the resulting subproblem. There are four cases to consider.

1.  $v$  and  $w$  are the rightmost (leftmost) roots of  $F'$  and  $G'$ , and  $S(F', G') = \text{right}$  (left).

Match  $v$  and  $w$  to get the desired subproblem.

2.  $v$  and  $w$  are the rightmost (leftmost) roots of  $F'$  and  $G'$ , and  $S(F', G') = \text{left}$  (right).

Note that at least one of  $F', G'$  is not a tree (since otherwise this is case (1)). Delete from one which is not a tree. After a finite number of such deletions we have reduced to case (1), either because  $S$  changes direction, or because both forests become trees whose roots are  $v, w$ .

3.  $v$  is the rightmost root of  $F'$ ,  $w$  is the leftmost root of  $G'$ .

If  $S(F', G') = \text{left}$ , delete from  $F'$ ; otherwise delete from  $G'$ . After a finite number of such deletions this reduces to one of the previous cases when one of the forests becomes a tree.


4.  $v$  is the leftmost root of  $F'$ ,  $w$  is the rightmost root of  $G'$ .

This case is symmetric to (3).  $\square$

We now turn to the  $\Omega(m^2n)$  lower bound on the number of relevant subproblems for any strategy.

**Lemma 4.** *For any strategy  $S$ , there exists a pair of trees  $(F, G)$  with sizes  $n, m$  respectively, such that the number of relevant subproblems is  $\Omega(m^2n)$ .*

*Proof.* Let  $S$  be an arbitrary strategy, and consider the trees  $F$  and  $G$  depicted in Fig. 4. According to lemma 3, every pair  $(F_v - v, G_w - w)$  where  $v \in F$  and  $w \in G$  is a relevant subproblem for  $S$ . Focus on such a subproblem where  $v$  and  $w$  are internal nodes of  $F$  and  $G$ . Denote  $v$ 's right child by  $v_r$  and  $w$ 's left child by  $w_\ell$ . Note that  $F_v - v$  is a forest whose rightmost root is the node  $v_r$ . Similarly,  $G_w - w$  is a



**Fig. 4.** The two trees used to prove an  $\Omega(m^2n)$  lower bound.

forest whose leftmost root is  $w_\ell$ . Starting from  $(F_v - v, G_w - w)$ , consider the computational path  $c_{v,w}$  that deletes from  $F$  whenever the strategy says left and deletes from  $G$  otherwise. In both cases, neither  $v_r$  nor  $w_\ell$  is deleted. Such deletions can be carried out so long as both forests are non-empty.

The length of this computational path is at least  $\min\{|F_v|, |G_w|\} - 1$ . Note that for each subproblem  $(F', G')$  along this computational path,  $v_r$  is the rightmost root of  $F'$  and  $w_\ell$  is the leftmost root of  $G'$ . It follows that for every two distinct pairs  $(v_1, w_1) \neq (v_2, w_2)$  of internal nodes in  $F$  and  $G$ , the relevant subproblems occurring along the computational paths  $c_{v_1, w_1}$  and  $c_{v_2, w_2}$  are disjoint. Since there are  $\frac{n}{2}$  and  $\frac{m}{2}$  internal nodes in  $F$  and  $G$  respectively, the total number of subproblems along the  $c_{v,w}$  computational paths is given by:

$$\sum_{(v,w) \text{ internal nodes}} \min\{|F_v|, |G_w|\} - 1 = \sum_{i=1}^{\frac{n}{2}} \sum_{j=1}^{\frac{m}{2}} \min\{2i, 2j\} = \Omega(m^2n)$$


$\square$

The  $\Omega(m^2n)$  lower bound established by Lemma 4 is tight if  $m = \Theta(n)$ , since in this case our algorithm achieves an  $O(n^3)$  runtime.

To establish a tight bound when  $m$  is not  $\Theta(n)$ , we use the following technique for counting relevant subproblems. We associate a subproblem consisting of subforests  $(F', G')$  with the unique pair of vertices  $(v, w)$  such that  $F_v, G_w$  are



the smallest trees containing  $F', G'$  respectively. For example, for nodes  $v$  and  $w$  with at least two children, the subproblem  $(F_v - v, G_w - w)$  is associated with the pair  $(v, w)$ . Note that all subproblems encountered in a computational path starting from  $(F_v - v, G_w - w)$  until the point where either forest becomes a tree are also associated with  $(v, w)$ .



**Fig. 5.** The two trees used to prove  $\Omega(m^2 n \log \frac{n}{m})$  lower bound.

**Lemma 5.** *For every strategy  $S$ , there exists a pair of trees  $(F, G)$  with sizes  $n \geq m$  such that the number of relevant subproblems is  $\Omega(m^2 n \log \frac{n}{m})$ .*

*Proof.* Consider the trees illustrated in Fig. 5. The  $n$ -sized tree  $F$  is a complete balanced binary tree, and  $G$  is a “zigzag” tree of size  $m$ . Let  $w$  be an internal node of  $G$  with a single node  $w_r$  as its right subtree and  $w_l$  as a left child. Denote  $m' = |G_w|$ . Let  $v$  be a node in  $F$  such that  $F_v$  is a tree of size  $n' + 1$  such that  $n' \geq 4m \geq 4m'$ . Denote  $v$ 's left and right children  $v_l$  and  $v_r$  respectively. Note that  $|F_{v_l}| = |F_{v_r}| = \frac{n'}{2}$ .

Let  $S$  be an arbitrary strategy. We aim to show that the total number of relevant subproblems associated with  $(v, w)$  or with  $(v, w_l)$  is at least  $\frac{n'm'}{4}$ . Let  $c$  be the computational path that always deletes from  $F$  (no matter whether  $S$  says left or right). We consider two complementary cases.

CASE 1:  $\frac{n'}{4}$  left deletions occur in the computational path  $c$ , and at the time of the  $\frac{n'}{4}$ th left deletion, there were fewer than  $\frac{n'}{4}$  right deletions.

We define a set of new computational paths  $\{c_j\}_{1 \leq j \leq \frac{n'}{4}}$  where  $c_j$  deletes from  $F$  up through the  $j$ th left deletion, and thereafter deletes from  $F$  whenever  $S$  says right and from  $G$  whenever  $S$  says left. At the time the  $j$ th left deletion occurs, at least  $\frac{n'}{4} \geq m' - 2$  nodes remain in  $F_{v_r}$ , and all  $m' - 2$  nodes are present in  $G_{w_l}$ . So on the next  $m' - 2$  steps along  $c_j$ , neither of the subtrees  $F_{v_r}$  and  $G_{w_l}$  is totally deleted. Thus, we get  $m' - 2$  distinct relevant subproblems associated with  $(v, w)$ . Notice that in each of these subproblems, the subtree  $F_{v_l}$  is missing exactly  $j$  nodes. So we see that, for different values of  $j \in [1, \frac{n'}{4}]$ , we get disjoint sets of  $m' - 2$  relevant subproblems. Summing over all  $j$ , we get  $\frac{n'}{4}(m' - 2)$  distinct relevant subproblems associated with  $(v, w)$ .

CASE 2:  $\frac{n'}{4}$  right deletions occur in the computational path  $c$ , and at the time of the  $\frac{n'}{4}$ th right deletion, there were fewer than  $\frac{n'}{4}$  left deletions.

We define a different set of computational paths  $\{\gamma_j\}_{1 \leq j \leq \frac{n'}{4}}$  where  $\gamma_j$  deletes from  $F$  up through the  $j$ th right deletion, and thereafter deletes from  $F$  whenever  $S$  says left and from  $G$  whenever  $S$  says right (i.e.,  $\gamma_j$  is  $c_j$  with the roles of left and right exchanged). Similarly as in case 1, for each  $j \in [1, \frac{n'}{4}]$  we get  $m' - 2$  distinct relevant subproblems in which  $F_{v_r}$  is missing exactly  $j$  nodes. All together, this gives  $\frac{n'}{4}(m' - 2)$  distinct subproblems. Note that since we never make left deletions from  $G$ , the left child of  $w_l$  is present in all of these subproblems. Hence, each subproblem is associated with either  $(v, w)$  or  $(v, w_l)$ .

In either case, we get  $\frac{n'}{4}(m' - 2)$  distinct relevant subproblems associated with  $(v, w)$  or  $(v, w_l)$ . To get a lower bound on the number of problems we sum over all pairs  $(v, w)$  with  $G_w$  being a tree whose right subtree is a single node, and  $|F_v| \geq 4m$ . There are  $\frac{m}{4}$  choices for  $w$  corresponding to tree sizes  $4j$  for  $j \in [1, \frac{m}{4}]$ . For  $v$ , we consider all nodes of  $F$  whose distance from a leaf is at least  $\log(4m)$ . For each such pair we count the subproblems associated with  $(v, w)$  and  $(v, w_l)$ . So the total number of relevant subprob-

lems counted in this way is

$$\begin{aligned}
\sum_{v,w} \frac{|F_v|}{4} (|G_w| - 2) &= \frac{1}{4} \sum_v |F_v| \sum_{j=1}^{\frac{m}{4}} (4j - 2) \\
&= \frac{1}{4} \sum_{i=\log 4m}^{\log n} \frac{n}{2^i} \cdot 2^i \sum_{j=1}^{\frac{m}{4}} (4j - 2) \\
&= \Omega(m^2 n \log \frac{n}{m})
\end{aligned}$$

□

**Lemma 6.** *For every strategy  $S$  and  $n \geq m$ , there exist trees  $F$  and  $G$  of sizes  $\Theta(n)$  and  $\Theta(m)$  for which  $S$  has  $\Omega(m^2 n (1 + \log \frac{n}{m}))$  relevant sub-problems.*

*Proof.* If  $m = \Theta(n)$  then this bound is  $\Omega(m^2 n)$  as shown in Lemma 4. Otherwise, this bound is  $\Omega(m^2 n \log \frac{n}{m})$  which was shown in Lemma 5. □

## 5 Conclusions

We presented a new  $O(n^3)$ -time and  $O(n^2)$ -space algorithm for computing the tree edit distance between two ordered trees. Our algorithm is not only faster than all previous algorithms in the worst case, but we have proved it is optimal within the broad class of decomposition strategy algorithms. As a consequence, any future improvements in terms of worst-case time complexity would have to find an entirely new approach. We obtain similar results when considering the sizes  $m$  and  $n$  of the input trees as separate parameters.

The novelty of our dynamic program is that it is both symmetric in its two inputs as well as adaptively dependant on them. This general notion may also be applied in other scenarios where the known dynamic programming solutions possess an inherent asymmetry.

The full version of the paper includes an explicit dynamic program for our algorithm, a proof of the  $O(n^2)$ -space complexity, and an adaptation of the algorithm for the edit distance problem on unrooted trees.

## References

1. P. Bille. Tree edit distance, alignment distance and inclusion. *Technical report TR-2003-23. IT University of Copenhagen*, 2003.
2. S. S. Chawathe. Comparing hierarchical data in external memory. In *Proceedings of the 25th International Conference on Very Large Data Bases*, pages 90–101, Edinburgh, Scotland, U.K., 1999.
3. W. Chen. New algorithm for ordered tree-to-tree correction problem. 40:135–158, 2001.
4. S. Dulucq and H. Touzet. Analysis of tree edit distance algorithms. In *Proceedings of the 14th annual symposium on Combinatorial Pattern Matching (CPM)*, pages 83–95, 2003.
5. D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM Journal of Computing*, 13(2):338–355, 1984.
6. P. N. Klein. Computing the edit-distance between unrooted ordered trees. In *Proceedings of the 6th annual European Symposium on Algorithms (ESA)*, pages 91–102, 1998.
7. P. N. Klein, S. Tirthapura, D. Sharvit, and B. B. Kimia. A tree-edit-distance algorithm for comparing simple, closed shapes. In *Proceedings of the 11th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 696–704, 2000.
8. P.B Moore. Structural motifs in RNA. *Annual review of biochemistry*, 68:287–300, 1999.
9. D. Shasha and K. Zhang. Simple fast algorithms for the editing distance between trees and related problems. *SIAM Journal of Computing*, 18(6):1245–1262, 1989.
10. K. Tai. The tree-to-tree correction problem. *Journal of the Association for Computing Machinery (JACM)*, 26(3):422–433, 1979.
11. G. Valiente. *Algorithms on Trees and Graphs*. Springer-Verlag, 2002.
12. R. A. Wagner and M. J. Fischer. The string-to-string correction problem. *Journal of the ACM*, 21(1):168–173, 1974.